(12) UK Patent Application (19) GB (11) 2 360 909 (13) A

(43) Date of A Publication 03.10.2001

(21) Application No 0007848.5

(22) Date of Filing 31.03.2000

(51) INT CL[7]
H04Q 7/22

(52) UK CL (Edition S)
H4L LDGP LRRMN L1H10 L201 L213

(56) Documents Cited
EP 0924897 A WO 95/16332 A US 5970412 A
US 5930242 A

(58) Field of Search
UK CL (Edition R) H4L LDGP LDH LDSWA LDSWC LDTX LFNA LFNC
INT CL[7] H04B 7/005, H04L 12/56, H04Q 3/00 7/22 7/36 7/38 11/04
ONLINE: WPI, PAJ, EPODOC

(71) Applicant(s)
Motorola Inc
(Incorporated in USA - Delaware)
Corporate Offices, 1303 East Algonquin Road, Schaumburg, Illinois 60196, United States of America

(72) Inventor(s)
Peter Jonathon Legg
Stephen Barrett
Stephen Hill

(74) Agent and/or Address for Service
Peter D Hudson
Motorola Limited, European Intellectual Property Department, Midpoint, Alencon Link, BASINGSTOKE, Hampshire, RG21 7PL, United Kingdom

(54) Abstract Title
A packet scheduler and method

(57) In a CDMA cellular communication network, packet data may be broadcast at a discontinuous rate using a packet scheduler (7) which utilises available communication resources in each cell more efficiently than for continuous transmission. To maintain quality of service requirements, each packet scheduler is associated a power bin (9,10,11) which represents the total power available for transmission and reception in the cell. If a base station is experiencing a higher level of traffic than its neighbours, the power bin for that cell may be increased with quality of service requirements being maintained by decreasing the power bins of neighbouring cells. Reducing the maximum transmission value for an under utilized base station, decreases intercell interference in neighbouring cells. This means that more packets can be scheduled to an overloaded cell (1) without increasing transmission power, whilst maintaining a target signal to interference ratio at a receiver. (4,5).

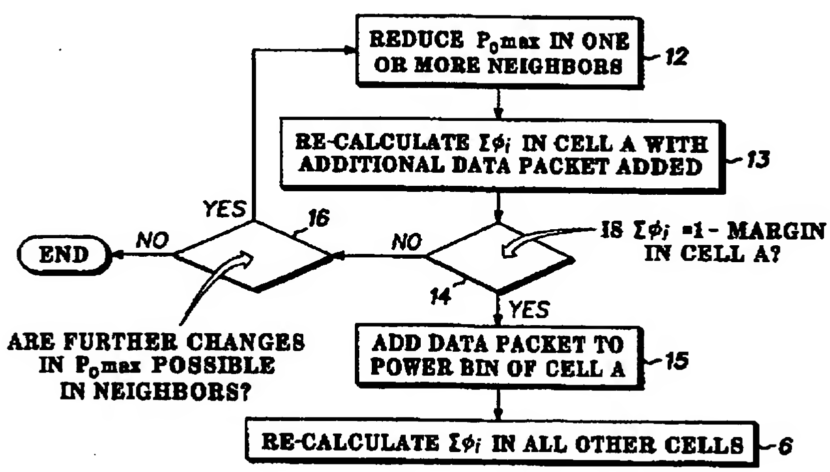

FIG. 5

GB 2 360 909 A

At least one drawing originally filed was informal and the print reproduced here is taken from a later filed formal copy.

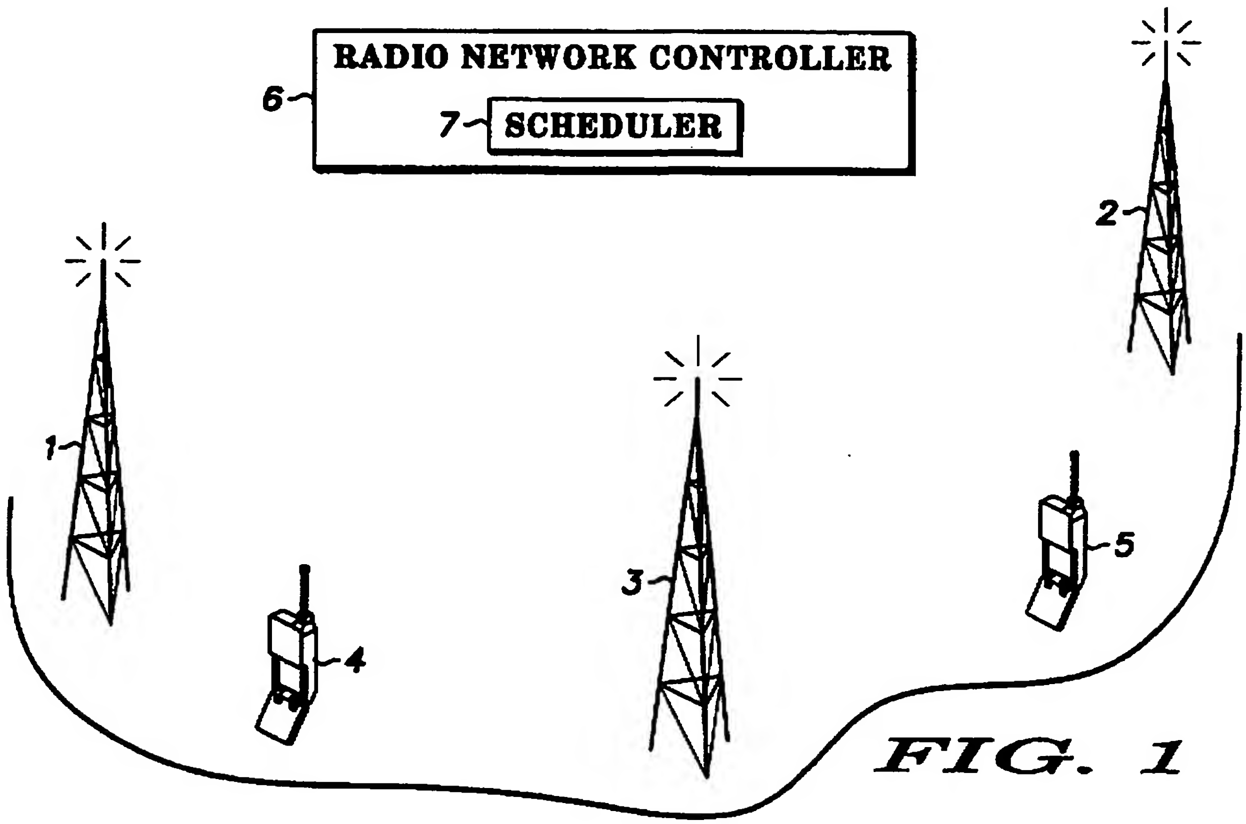
RADIO NETWORK CONTROLLER
6
7
SCHEDULER
1
2
3
4
5

FIG. 1

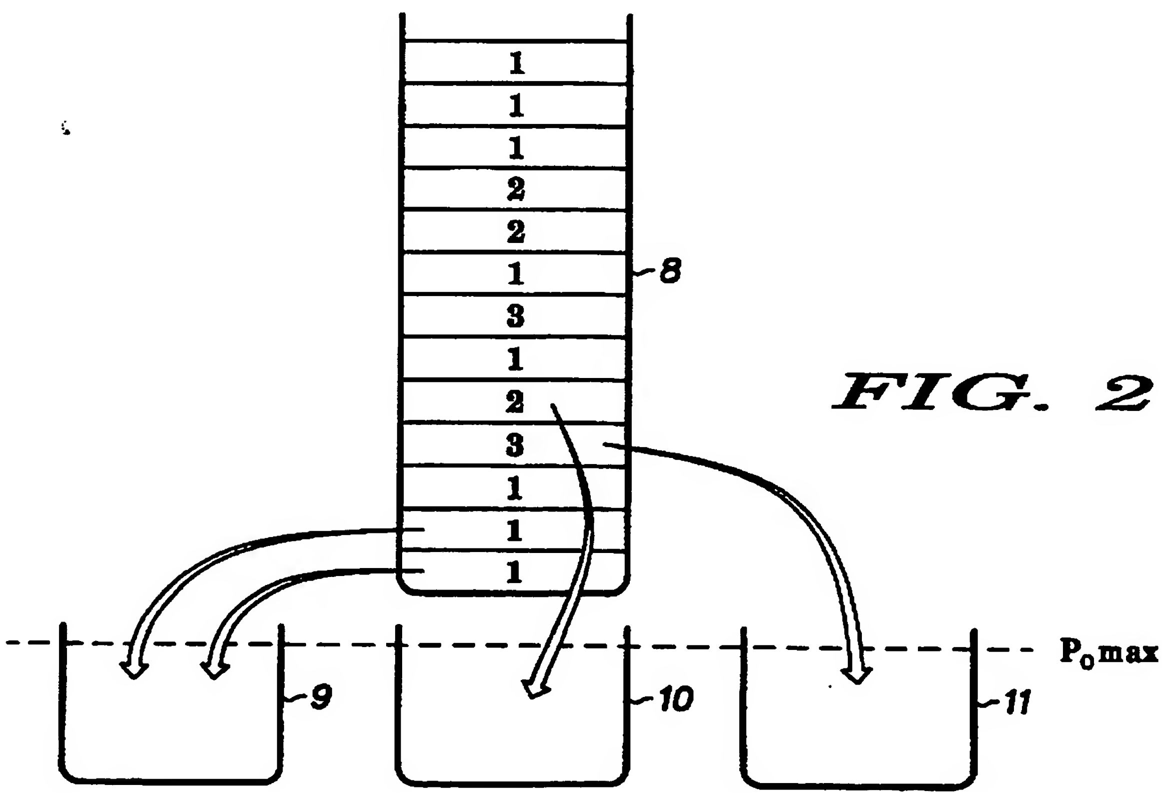
1
1
1
2
2
1
3
1
2
3
1
1
1
8
P_0max
9
10
11

FIG. 2

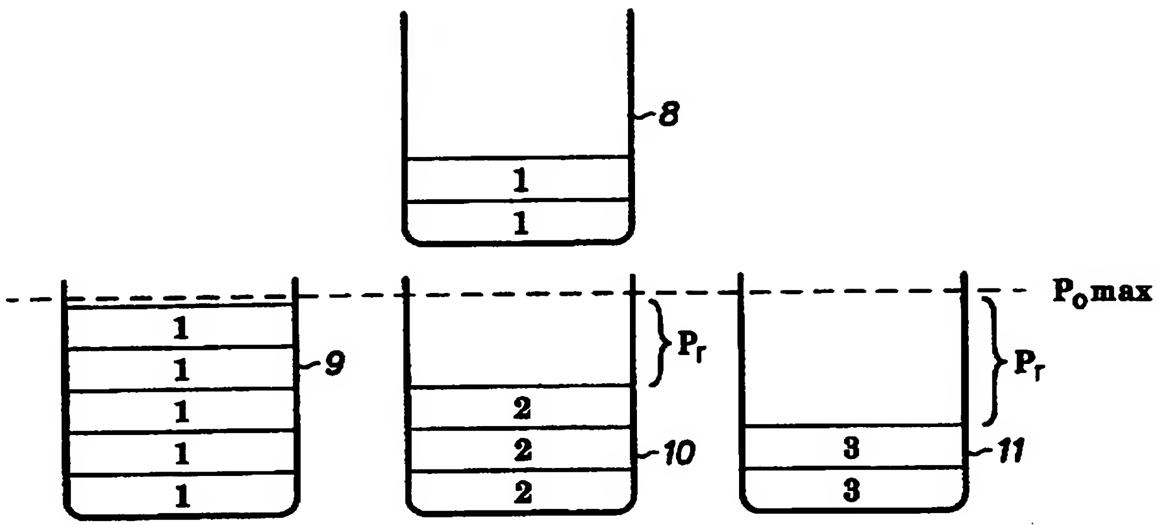

8
1
1
$P_o$max
1
1
1
1
1
9
P_r
2
2
2
10
P_r
3
3
11

FIG. 3

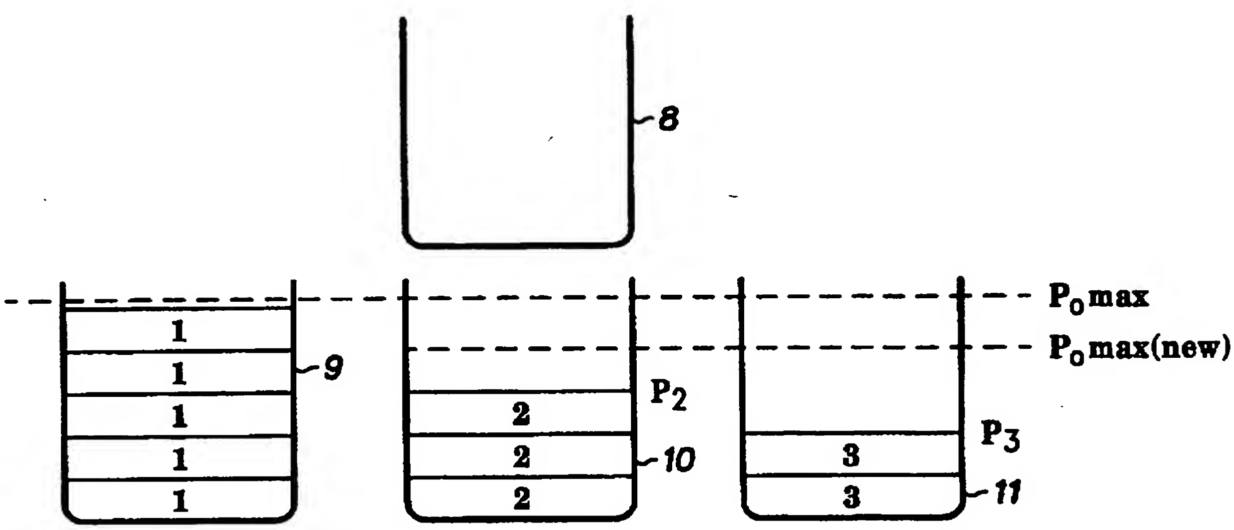

8
$P_o$max
$P_o$max(new)
1
1
1
1
1
9
P_2
2
2
2
10
P_3
3
3
11

FIG. 4

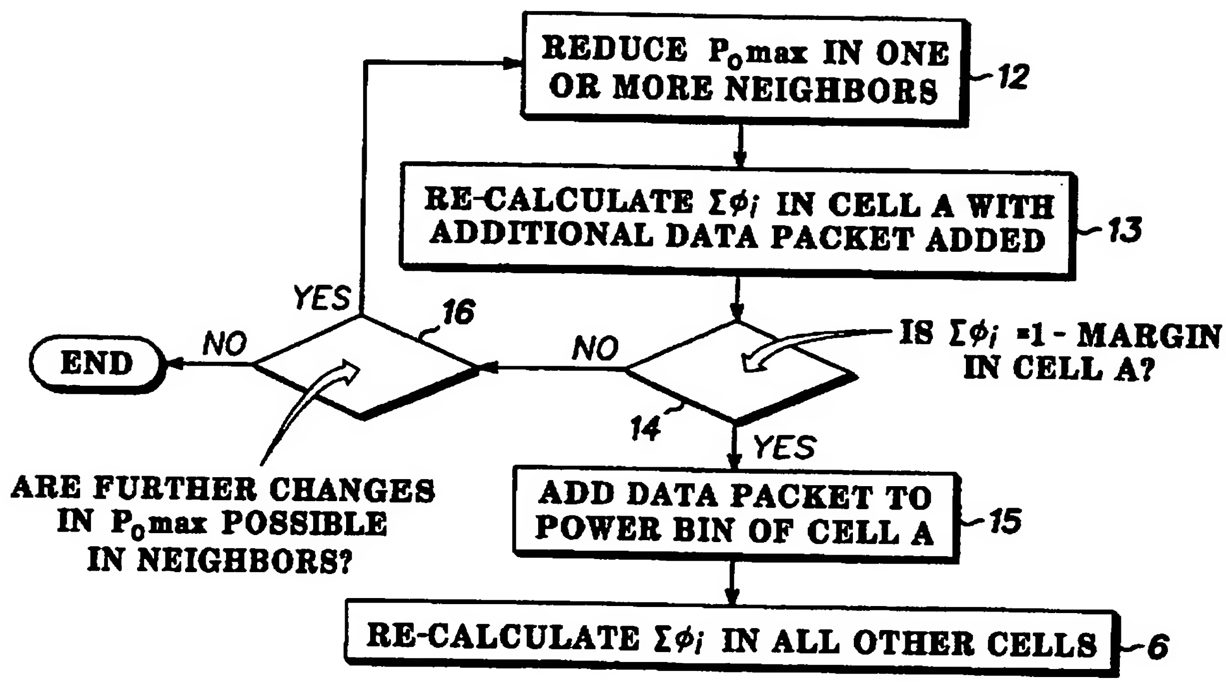
REDUCE P_0max IN ONE OR MORE NEIGHBORS
12
RE-CALCULATE $\Sigma\phi_i$ IN CELL A WITH ADDITIONAL DATA PACKET ADDED
13
IS $\Sigma\phi_i$ =1 - MARGIN IN CELL A?
14
NO
YES
ADD DATA PACKET TO POWER BIN OF CELL A
15
RE-CALCULATE $\Sigma\phi_i$ IN ALL OTHER CELLS
6
ARE FURTHER CHANGES IN P_0max POSSIBLE IN NEIGHBORS?
16
YES
NO
END

FIG. 5

2360909

A PACKET SCHEDULER AND METHOD

Field of the Invention

This invention relates to a packet scheduler and method therefor, and in particular to a packet scheduler for a cellular communication system for mobile communication.

Background of the Invention

In a cellular communication system remote terminals (typically mobile stations) communicate with a fixed base station. Communication from the mobile stations to the base station is known as uplink and communication from the base station to the mobile stations is known as downlink. The total coverage area of the system is divided into a number of separate cells, each covered by a single base station. The cells are typically geographically distinct with an overlapping coverage area with neighbouring cells.

As a mobile station moves from the coverage area of one cell to the coverage area of another cell, the communication link will change from being between the mobile station and the base station of the first cell, to being between the mobile station and the base station of the second cell. This is known as a handover or handoff.

All base stations are interconnected by a fixed network. This fixed network comprises communication lines, switches, interfaces to other communication networks and various controllers required for operating the network. A call from a mobile station is routed through the fixed network to the destination specific for this call. If the call is between two mobile stations of the same communication system the call will be routed through the fixed network to the base station of the cell in which the other mobile station currently is. A connection is thus established between the two serving cells

through the fixed network. Alternatively, if the call is between a mobile station and a telephone connected to the Public Switched Telephone Network (PSTN) the call is routed from the serving base station to the interface between the cellular mobile communication system and the PSTN. It is then routed from the interface to the telephone by the PSTN.

A cellular mobile communication system is allocated a frequency spectrum for the radio communication between the mobile stations and the base stations. This spectrum must be shared between all mobile stations simultaneously using the system.

One method of sharing this spectrum is by a technique known as Code Division Multiple Access (CDMA). In a Direct Sequence CDMA (DS-CDMA) communication system, the signals are prior to being transmitted, multiplied by a high rate code whereby the signal is spread over a larger frequency spectrum. A narrowband signal is thus spread and transmitted as a wideband signal. At the receiver the original narrowband signal is regenerated by multiplication of the received signal with the same code. A signal spread by use of a different code will at the receiver not be de-spread but will remain a wide band signal. In the receiver the majority of interference caused by interfering signals received in the same frequency spectrum as the wanted signal can thus be removed by filtering. Consequently a plurality of mobile stations can be accommodated in the same wideband spectrum by allocating different codes for different remote terminals. Codes are chosen to minimise the interference caused between mobile stations typically by choosing orthogonal codes when possible. A further description of CDMA communication systems can be found in 'Spread Spectrum CDMA Systems for Wireless Communications', Glisic & Vucetic, Artech house Publishers, 1997, ISBN 0-89006-858-5. Examples of CDMA cellular communication systems are IS 95 standardised in North America and the Universal Mobile Telecommunication System (UMTS) currently under standardisation in Europe.

Traditional traffic in mobile cellular communication systems has been circuit switched voice data where a permanent link is set up between the communicating parties. In the future it is envisaged that data communication will increase substantially and typically the requirements for a remote terminal to transmit data will not be continuous but will be at irregular intervals. Consequently it is inefficient to have a continuous link setup between users and instead a significant increase in packet based data traffic is expected, where the transmitting remote terminal seeks to transmit the data in discrete data packets when necessary. An example of a packet based system is General Packet Radio Service (GPRS) introduced to the Global System for Mobile communication (GSM). Further details on data packet systems can be found in 'Understanding data communications: from fundamentals to networking, 2nd ed.', John Wiley publishers, author Gilbert Held, 1997, ISBN 0-471-96820-X.

In a packet based system where a high number of mobile stations may require resources for packet transmissions at unknown and irregular intervals it is important for optimal utilisation of the limited resource to schedule the order and time for transmission of the individual packets. This becomes even more important when different data packets have different requirements with respect to delay, bit error rate etc. Therefore most packet based systems contain schedulers which control when the individual data packets are transmitted and therefore share the available resource, whether time-slots in a TDMA system or power and codes in a CDMA system. An introduction to schedulers can be found in 'Service discipline for guaranteed performance service in packet-switching networks', Hui Zhang, Proceedings of the IEEE, volume 83, no. 10, October 1995.

However, known schedulers have been optimised for different environments than CDMA systems. For example, scheduling algorithms used for GPRS are optimised for a Time Division Multiple Access (TDMA) system and therefore not optimal for CDMA systems where codes and power must be shared.

This invention aims to provide a packet scheduler suitable for use in a CDMA system and which enables an efficient sharing of network resource.

Summary of the Invention

According to a first aspect, the present invention comprises apparatus for scheduling queued data packets for transmission between a plurality of base stations, each serving a respective cell in a cellular communications network, and a plurality of mobile stations located therein, the apparatus including,

means for assigning an associated power bin to each cell into which at least some of the queued data packets are to be placed,

means for setting a maximum power level value for each power bin,

means for identifying a first power bin associated with an overloaded cell and second power bin associated with an under-loaded cell,

means for reducing the maximum power level of the second power bin, thereby reducing intercell interference in the overloaded cell,

and means for allocating further data packets from the queue to the first power bin.

According to a second aspect, the present invention comprises a method for scheduling queued data packets for transmission between a plurality of base stations, each serving a respective cell in a cellular communications network, and a plurality of

mobile stations located therein, the method including the steps of,

assigning an associated power bin to each cell into which at least some of the queued data packets are to be placed,

setting a maximum power level value for each power bin,

identifying a first power bin associated with an overloaded cell and second power bin associated with an under-loaded cell,

reducing the maximum power level ~~of the~~ second power bin, thereby reducing intercell interference in the overloaded cell,

and allocating further data packets from the queue to the first power bin.

In one embodiment, a maximum power level value ($P_o$max) is set for each power bin.

A total power level of each power bin resulting from the data packets which are destined therefor is calculated during a scheduling period, the calculation taking account of intercell interference based on transmissions from each base station being set at a level of $P_o$max. A first power bin whose total power level is substantially equal to $P_o$max is identified and a second power bin whose total power level is less than $P_o$max is also identified. The maximum power level of the second power bin is set to a new, lower value $P_o$max (new).

A re-calculation of the total power level of the first power bin is performed in the cell based on downlink transmissions associated with the second power bin being set at a level of $P_o$max (new). Subsequently further data packets can be allocated from the queue to the first power bin. In an alternative

embodiment, the calculations and re-calculations of total power levels of power bins are performed based on uplink transmissions from mobile stations operating within the communications system.

Thus the invention allows re-distribution of network resource from an under-utilised base station to an overloaded one.

In a preferred embodiment the apparatus comprises a scheduler which is located within a radio network controller and all the steps of the method are performed by the scheduler.

Brief Description of the Drawings

Some embodiments of the invention will now be described, by way of example only, with reference to the drawings of which:

Fig.1 is a schematic block diagram of a cellular communications system operating in accordance with the invention;

Figs. 2, 3 & 4 are schematic diagrams illustrating the operation of a packet scheduler in accordance with the invention;

And Fig. 5 is a flow chart illustrating operation of a specific embodiment of the invention.

Detailed Description of the Preferred Embodiments

The following embodiments are described within the context of, the current approach for the standardisation of UMTS but it will be apparent that the invention is not limited to this application.

FIG 1 shows a schematic diagram of an embodiment of a CDMA communication system in accordance with an embodiment of the invention. The communication system has a number of base

stations 1, 2, 3 each covering a geographical area and thereby defining a cell. A number of mobile stations 4, 5 are associated with the communication system and communicate to each other or to other systems via the base stations 1, 2, 3.

The base stations are connected to a common Radio Network Controller (RNC) 6. The RNC further provides gateways to other communication systems such as the fixed public telephone system (not shown) and contains a scheduler 7.

Each of the mobile stations 4, 5 have independent communication needs and communicate by use of data packets. The mobile stations 4, 5 may require different services and can for example be Internet browsers, telephones or data terminals. Each remote terminal may also request different services at different times.

The resource requirement for each individual mobile station may vary significantly over time so that a mobile station may sometimes require no transmissions and at other times require long transmissions at high data rate. The resource requirement for each mobile station from the communication network can thus vary significantly and in order to ensure that the available network capacity is used optimally an efficient scheduling of packets for the different mobile stations is required. This task is performed by the scheduler 7.

In a frame based communication system such as UMTS the communication is divided into discrete time intervals or frames and the communication resource is allocated on a per frame basis. In UMTS, packets to be transmitted are scheduled during one time frame and transmitted during a subsequent frame.

In CDMA systems the total data throughput is affected by the interference at a mobile station's receiver caused by transmissions from base stations other than its serving base station. A successful sharing of the available resources of the system requires a knowledge of these interference levels so that

the signal to noise ratio for each connection is appropriate for the desired quality of service (QoS). In order to perform an optimal allocation, transmissions from non-serving base stations must therefore be taken into account and scheduling in all cells must be done simultaneously.

Consider sharing the power in a single cell. The scheduler's task is to divide the available transmit power of the base station (BS) between data packets removed in sequence from the queue. For each transmission the signal to noise ratio (SNR) at the destination mobile station must equal that necessary to achieve the agreed QoS. The SNR at mobile station *i* is given by the following equation:

$$SNR_i = \frac{P_0\phi_i / L_i}{\alpha P_0(1-\phi_i)/L_i + I_{\text{intercell},i} + P_{thermal}} \qquad \text{...(1)}$$

P_0 = BS transmit power
ϕ_i = fraction of power assigned to user i
L_i = path loss from BS to user i
α = loss of orthogonality factor
$I_{\text{intercell},i}$ = intercell interference at user i
P_{thermal} = thermal noise power added before the receiver

To determine the transmit power to a mobile, equation (1) must be solved for φi.

$$\phi_i = \frac{SNR_i\left(P_{th}L_i + I_{\text{intercell},i}L_i + \alpha P_0\right)}{P_0\left(1+SNR_i\alpha\right)} \qquad \text{...(2)}$$

Φ*i* Is equivalent to one or more data packets

The path loss to a mobile station, the loss of orthogonality factor and the thermal noise power can all be measured by known techniques and reported to the scheduler. P_o can be set by the network operator and so can a target SNR. Intercell interference can be calculated from knowledge of transmit power of the other base stations in the network and from the path loss.

Thus, by substituting the relevant values into equation (2), a value for fractional power of every mobile station operative in a cell can be found. The total power transmitted (during a frame) from a base station will then be the sum of the individual Φ_i values. ($\Sigma\Phi_i$) multiplied by P_o.

However, the total transmit power of any base station is limited for practical reasons (eg. Depending on the capabilities of the power amplifiers). Say that this limit is the same for all base stations and has a value $P_o$max. Hence, the number of packets that can be transmitted by a particular base station during a frame is similarly limited.

It will be appreciated, however, from inspection of equation (2) that if the intercell interference (for a particular cell) should decrease, then for a fixed target SNR, Φ_i will also decrease and hence, so will $\Sigma\Phi_i$. This means that more packets can be scheduled to and transmitted from the base station serving that cell without exceeding $P_o$max.

Fig. 2 shows a schematic representation of the constituent parts of the scheduler 7. A single queue of packets 8 is allocated as necessary to one of three power bins 9, 10, 11 which relate respectively to base stations 1, 2 and 3. Packets labelled 1 are destined for transmission by base station 1, packets 2 are

for base station 2 and packets 3 are for base station 3. A maximum value for the power bins 9, 10, 11 is set at $P_o$max.

A power bin represents the total available power at a base station, for either reception or transmission, which is to be shared by i users. The height of a data packet shown in the power bins of Figs. 2, 3 and 4 represents the amount of power required by a user per frame/scheduling period.

Fig. 3 shows the situation part-way through a frame/scheduling period. The scheduler has scheduled and calculated the values for Φ_i for a number of packets destined for base stations 1, 2 3 and begun to fill up the power bins 9, 10, 11 accordingly. In order to calculate a value for intercell interference, the value for power transmitted by every base station is taken to be $P_o$max. In Fig. 3 the power bin 9 for base station 1 is completely full with some packets in the queue 8 still remaining to be scheduled. On the other hand, power bins 10 and 11 are not full yet no more packets are required to be scheduled for base stations 2 and 3 during this frame. Thus base stations 2 and 3 have spare resource which could be utilised by the over-loaded base station 1.

Without the benefit of the present invention, when a power bin of a certain base station is filled (it reaches $P_o$max) then further packets directed this base station must wait until the next scheduling period/frame.

At the end of the scheduling, base stations 2 and 3 will transmit at a lower power because there are not sufficient packets for the cells to reach $P_o$max (no packets queued for the cells remain in the single queue.) Therefore, there is some unused capacity in the network.

Furthermore, as packets are skipped once the power bin for a base station is filled, coverage of hot spots (which occur when there is a lot of high priority traffic in a single cell) is prevented.

The present invention can utilise the unused capacity and further accommodate hot spots as follows.

$P_o$max for power bins 10 and 11 are lowered to a new value $P_o$max (new) [chosen by one of several methods to be discussed below]. The scheduler 7 then re-calculates the values for Φ_i for the base station 1 based on values for intercell interference determined from the adjusted values $P_o$max (new). The sum of the new Φ_i values will now have a value of less than unity. This means that more packets (remaining in queue 8) can be scheduled to base station 1 into power bin 9) yet still keeping total transmit power within the $P_o$max limit.

The re-configuration of the power bins at the end of the re-calculation procedure is shown in Fig. 4. At the end of the frame, base station 1 transmits all the queued packets for its cell at a total power of $P_o$max. Base stations 2 and 3 transit at a total power level of P_2 and P_3 both being less than $P_o$max (new). In Fig. 4, all packets in queue 8 have been successfully scheduled during the one frame.

Hence to accommodate hot spots, the invention permits the borrowing of downlink capacity from neighbouring cells.

One option for selecting which cells should have their maximum power reduced, is to choose those cells which generate the greatest interference to mobile stations lying within the hot spot cell.

In a further embodiment, the scheduler can be configured to redistribute any unused capacity within the network based on a re-calculation of the Φ_i values (and of $\Sigma\Phi_i$) in one or more neighbouring cells or all cells in the network.

A hot-spot is defined as a cell in which there is high priority traffic queued which cannot all be serviced if the network capacity is shared equally amongst all cells ($P_o$max equal in all cells). Hot-spots have limited lifespans, otherwise the network operator would deploy greater infrastructure in these areas to handle the greater demand. Examples include cells ~~around~~ a football ground, this cell becomes 'hot' at half-time and the end of play as people phone home; a stretch of motorway when there has been an accident and a queue of cars develops. It is of considerable benefit to an operator if hot-spots can be handled efficiently by careful radio resource management over a number of cells.

To handle a hot-spot it is necessary to be able to allocate packets from the single queue 8 which, conventionally is ordered in terms of transmission priority, *without skipping packets*. This is possible by borrowing capacity from other cells, in the manner described above with reference to Fig. 4.

In order to choose the appropriate values for $P_o$max (new) for any underused base station or base stations, the scheduler may keep reducing $P_o$max of the appropriate base station(s) and carry out the re-calculation of Φ_i values procedure until enough space has been created in the over-subscribed power bin to accommodate the queued packets. It is preferable to put a lower limit on $P_o$max (new) so that the capacity of the reduced power bins is not disadvantaged, ie there should be enough resource left available in the reduced bins so that their own packets which are already scheduled can be transmitted.

Different strategies may be employed in determining how to arrive at the best $P_o$max (new) value. For example, it is possible to reduce $P_o$max in cells for which the path loss from the base station to the mobile stations in the overloaded cell is smallest. This will require a smaller reduction than for cells which are 'further' (in the path loss sense) from the overloaded cell. In other words, it is best to borrow capacity from the neighbouring cells.

Alternatively, it is possible to reduce $P_o$max in proportion to the remaining available transmit power (designated Pr in Fig. 3) in the cells which have the greatest spare capacity.

As an alternative to this technique, the buffered traffic in queue 8 for each cell may be determined by the scheduler 7 prior to filling the power bins. Using this information the scheduler 7 can more accurately determine which cells will have spare capacity and therefore should have their $P_o$max re-defined.

Making the changes to a minimum number of base station $P_o$max values minimises complexity of the scheduling process. The process is simpler if only $P_o$max in one cell needs to be changed.

Using equation (2) the scheduler 7 forms an equation for the sum of φi (including an additional packet [s]) in the overloaded cell, and solves this for a $P_o$max (new) value of the single cell whose maximum power has been adjusted. If the scheduler were to adjust the $P_o$max in a number of cells simultaneously the mathematics is more complex and an iterative approach as outlined in Figure 5 (see below) is more appropriate. Following the accommodation of the additional packet(s) the scheduler recalculates the φi values allocated in cells other than over-loaded cell (since the intercell interference may have changed).

Alternatively, it may be less complex to recalculate the φi values for packets in a cell only when the scheduler is arranging for the addition of a packet to that same cell. It is also possible to calculate the new φi values for all packets in the one cell in one calculation.

Optionally, a safety margin may be introduced so that the sum of the desired recalculated φi values is $\Sigma\varphi i$: = 1-margin.

Figure 5 shows a flow chart illustrating the adjustment of the maximum transmit power ($P_o$max) in one or more cells to accommodate the allocation of an additional packet to a cell "A" which is a "hot spot". All steps are performed in the scheduler 7.

In Fig. 5, in step 12, $P_o$max is reduced for one or more neighbouring base stations. In step 13, the scheduler 7 re-calculates $\Sigma\Phi_i$ including an additional data packet in cell A. In step 14 it tests the value of $\Sigma\Phi_i$ for cell A to see if it lies within a pre-defined margin of value 1. If it does, then the scheduler 7 adds the additional data packet to cell A's power bin (step 15). (It may also go on to calculate Φ_i values in all other cells [step 6]). If not, then it checks to see if further changes in $P_o$max for neighbours is possible (step 16). If so, the process repeats. If not, the process ends.

Employing the method for dealing with hot spots as described above has a limitation in that the overall capacity of the system can be less than if the capacity is shared equally amongst all cells. This compromise may be addressed by considering a hybrid method which employs a priority obedience factor, gamma. In such a hybrid method the extent to which maximum transmit power may be reduced below the nominal value of Pmax is controlled by gamma. When gamma is set to 1, this is the

hot spot method as described above and the transmit power $P_o$max (new) may be set to zero if necessary. In addition, gamma values of between zero and 1 can be used where gamma controls either the maximum reduction of the transmit power below Pmax (i.e. this is a hard limit), or aspects of a probability density function for the transmit power reduction, for example, the standard deviation (i.e. this is a soft limit). In this configuration there is both an attempt to follow the strict priority order and also to achieve good utilisation of the network.

In an alternative embodiment with reference again to Fig. 4, the $P_o$max (new) values of the under-utilised power bins (eg bins 10 and 11) are set at P_2 and P_3 respectively. The recalculation of Φ_i values and further allocation of queued packets to power bin 9 is then performed as above. In general terms, in this embodiment which permits sharing of capacity among cells, the scheduler freezes the value of $P_o$max in a power bin which is unfilled, to the level reached at some pre-determined point during the frame and then considers re-allocation of the further queued packets to cells whose power bins are already full. A procedure for this embodiment is as follows. At the pre-determined point in a frame, the scheduler sets $P_o$max (new) for each power bin to the filled power level in each bin. For power bins thus having $P_o$max (new) less than $P_o$max, then these levels are effectively "frozen". Next, the scheduler 7 re-calculates Φ_i values in cells which have $P_o$max (new) equal to $P_o$max (ie full bins). The scheduler then continues scheduling and adds more packets where possible to the power bins not "frozen".

While the preferred embodiments have been described with respect to scheduling data packets on a downlink, it is to be understood that similar principles can be employed on an uplink, thus allowing system capacity sharing between cells.

In this case, a power bin now represents received power at a base station, and this includes both useful power and intercell interference. The total uplink power resource in a cell (P_{total}) is a constant based on the characteristics of the UMTS uplink. The capacity of the uplink is limited by the need to meet the signal to noise ratio for every transmission (SNR). P_{total} is set to such a level to ensure a high sustainable throughput.

At the beginning of each frame, a maximum useable power for resource scheduling in any cell (P_{max}) is set to an appropriate level based on the mean ratio of the received power at a base station from the mobile station it serves to the received power from all other mobile stations. (This ratio is set a approximately 1:0.6). This ensures that there is adequate headroom within any cell to allow for the neighbouring cells to be fully loaded and thus use their power resource fully.

The invention as described above in the downlink case can be similarly applied in the uplink case in order to transfer power resource between cells whilst maintaining the Signal to Noise Ratio (SNR) and therefore the desired QoS. Again, the principle employed here is to reduce the P_{max} in cells neighbouring a highly loaded cell. This means that while P_{total} remains constant, the split between usable uplink power P_{max} at the receiver of a highly loaded base station and intercell interference is altered by lowering the amount of allowed intercell interference. This reduces the interference expected in the highly loaded cell, allowing P_{max} to be increased, and thus additional packets to be scheduled.

In a preferred embodiment, if power resource is transferred from one cell to another, then extra interference will be generated and will impact on the SNR at the cells which power has been taken from. Sufficient headroom (60% of the additional power requirement) must be available at all adjacent cells to enable the extra interference to be handled without reducing the SNR of

already scheduled packets. Thus P_{max} must be reduced by say 60% of the additional power requirement at all adjacent cells.

P_{max} can now be further reduced at adjacent cells to create the additional headroom for scheduling in the hot-spot cell. It is assumed that the interference from any cell is distributed equally between all neighbouring cells. The reduction of P_{max} by 1 unit in an adjacent cell will result in a reduction in interference and therefore an an available increase in P_{max} by 0.6 /N in the hot spot cell, where N is the number of cells neighbouring the hot spot cell. The preferred embodiment implements the distribution of P_{max} reduction by further reducing P_{max} in at least one of the adjacent cells (for example, the most lightly loaded neighbour) until the required P_{max} increase in the hot spot cell is reached. Alternatively, P_{max} of one adjacent cell is reduced until the level of another adjacent cell is reached, then both are reduced at the same rate until the required P_{max} increase in the hot spot cell is reached.

When the above conditions are met and the required extra headroom is found, additional packet(s) can be scheduled; otherwise they must be skipped.

CLAIMS

1. Apparatus for scheduling queued data packets for transmission between a plurality of base stations, each serving a respective cell in a cellular communications network, and a plurality of mobile stations located therein, the apparatus including,

means for assigning an associated power bin to each cell into which at least some of the queued data packets are to be placed,

means for setting a maximum power level value for each power bin,

means for idendifying a first power bin associated with an overloaded cell and second power bin associated with an under-loaded cell,

means for reducing the maximum power level of the second power bin, thereby reducing intercell interference in the overloaded cell,

and means for allocating further data packets from the queue to the first power bin.

2. A method for scheduling queued data packets for transmission between a plurality of base stations, each serving a respective cell in a cellular communications system, and a plurality of mobile stations located therein, the method including the steps of,

(a) assigning an associated power bin to each cell into which at least some of the queued data packets are to be placed,

(b) setting a maximum power level value for each power bin,

(c) identifying a first power bin associated with an overloaded cell and second power bin associated with an under-loaded cell,

(d) reducing the maximum power level of the second power bin, thereby reducing intercell interference in the overloaded cell,

(e) and allocating further data packets from the queue to the first power bin.

3. A method according to claim 2 in which step (b) includes the step of setting a maximum power level value $P_o max$ for each power bin; further including the step of calculating during a scheduling period, a total power level of each power bin resulting from the data packets which are destined therefor, the calculation taking account of intercell interference based on transmissions from each base station being set at a level of $P_o max$, and in which step (c) includes the step of identifying the first power bin as one whose total power level is substantially equal to $P_o max$ and the second power bin and one whose total and power level is less than $P_o max$; and in which step (d) includes the step of resetting the maximum power level of the second power bin to a new, lower value $P_o max$ (new), and including the further step of re-calculating the total power level of the first power bin based on downlink transmissions in the case associated with the second power bin being set at a level of $P_o max$ (new).

4. A method according to claim 3 in which the cell associated with the second power bin is one which contributes a high level of interference, compared with other cells in the system, to mobile stations located in the cell associated with the first power bin.

5. A method according to either of claims 3 or 4 and including the further step of repeating the steps of resetting, re-calculating and allocating until a pre-determined limit for P_o (max) new is reached.

6. A method according to any of claims 3 to 5 including the further step of re-calculating the total power level of the power bin associated with at least one cell in the system other than the first cell after the re-setting step.

7. A method according to any of claims 3 to 6 in which $P_o max$ (new) is set to a level dependent upon the magnitude of the difference between $P_o max$ and the total power level of the second power bin.

8. A method according to any of claims 4 to 6 in which $P_o max$ (new) is set to the total power level for the second power bin.

9. A method according to claim 2 in which step (b) includes the step of setting a maximum power level value Pmax for each power bin where Pmax is set to a level based on a signal to noise ratio of a signal received at each base station, and in which step (d) includes the steps of reducing Pmax in proportion to the degree of overload of the overloaded cell in cells adjacent to the overloaded cell and further reducing Pmax of the power bin of at least one adjacent cell.

10. A computer program product comprising a medium on or in which is recorded a program which, when executed in a computer-controlled system, will perform the method of any of claims 2-9.

11. A data packet scheduler substantially as hereinbefore described with reference to the drawings.

12. A method for scheduling queued data packets substantially as described with reference to the drawings.

The Patent Office

INVESTOR IN PEOPLE

Application No: GB 0007848.5
Claims searched: 1-12

Examiner: Robert Macdonald
Date of search: 29 September 2000

Patents Act 1977
Search Report under Section 17

Databases searched:

UK Patent Office collections, including GB, EP, WO & US patent specifications, in:

UK Cl (Ed.R): H4L(LDGP, LDH, LDSWA, LDSWD, LDTX, LFNA, LFNC)

Int Cl (Ed.7): H04B(7/005); H04L(12/56); H04Q(3/00, 7/22, 7/36, 7/36, 7/38, 11/04)

Other: ONLINE: WPI, PAJ, EPODOC

Documents considered to be relevant:

Category	Identity of document and relevant passage		Relevant to claims
A	EP 0924897 A	(AT+T) See paragraphs 12,17,19, and 25	
A	WO 95/16332 A	(BRITISH TELECOM.) See whole document.	
A	US 5970412	(NICHOLAS FRANK MAXEMCHUK) See whole document	
A	US 5930242	(NEC) See whole document	

X	Document indicating lack of novelty or inventive step	A	Document indicating technological background and/or state of the art.
Y	Document indicating lack of inventive step if combined with one or more other documents of same category.	P	Document published on or after the declared priority date but before the filing date of this invention.
		E	Patent document published on or after, but with priority date earlier than, the filing date of this application.
&	Member of the same patent family		

This Page Blank (uspto)